Tomographic Multi-Target Tracking and Its Relationship to PHD Filters

Lecture 2

Roy L. Streit
Metron, Inc.
streit@metsci.com

Department of Electrical and Computer Engineering
University of Massachusetts – Dartmouth
r.streit@ieee.org

NATO Lecture Series, SET-157
Multisensor Fusion: Advanced Methodologies and Applications

6-7 June 2011, Porto (PT)
9-10 June 2011, Paris (FR)
20 – 21 June 2011Lancaster (UK)
23 – 25 June 2011, Ankara (TR)

Topics – Lecture 2

- More examples
  - Kinematic data augmented with feature data
    - "Boosts" the underlying likelihood function
  - Track extraction by heritage method
  - Examples / Movies on simulated multistatic sonar data
- Method 3: The first moment approximation
  - Mean field approximation and scaling
  - First moment and the Janossy density

  } *Two methods but really only one*
  - Janossy density = the symmetrized Bayes information update
- Other topics
  - Regularization
  - Bayesian Cramer-Rao Bounds (BCRB)
  - Comparison to LRT (Likelihood Ratio Tracking)
- SPECT and the Miller-Snyder-Miller algorithm
- Multiple Sensor Multi-Target Intensity Filtering
  - Miller-Snyder-Miller algorithm
- Variance reduction via multisensor filtering
- Concluding remarks

Feature Data

- Two hypotheses about the data point z
  - H0 = z corresponds to clutter
  - H1 = z corresponds to real target
- Augment kinematic measurement z with feature vector u
- Use Bayes Theorem

  p(z,u|x,H1) = p(u|z,x) p(z|x)

  p(z,u|x,H0) = p(u|,z,phi) p(z|phi)
- Assume these four pdfs are known
  - "Fantasy features" are features with a priori specified pdfs
  - Estimating them accurately is important in practice
- Information update factor of PHD filter with feature data:

$$L_k(Z,U \mid x) = 1 - P_k^D(x) + \sum_{(z,u)\in(Z,U)=Data\ Set} \frac{\frac{p_k(u|z,x)}{p_k(u|z,\varphi)}\, p_k(z \mid x) P_k^D(x)}{\lambda_k(z \mid \varphi) + \int_S \frac{p_k(u|z,x)}{p_k(u|z,\varphi)}\, p_k(z \mid s) P_k^D(s)\, f_{k|k-1}(s)\, ds}$$

- Likelihood ratio of the feature data u
  - "Boosts" the likelihood of the kinematic data z
- Reduces to usual case when no feature data are available

Implementing the PHD and iFilters

- Two basic models of the posterior intensity
  - Gaussian sum
  - Sequential Monte Carlo (SMC), or "Particles"
- Particles are unevenly distributed over the state space
  - Represent posterior target intensity at current time
- Basic steps in particle filter from time t(k-1) to time t(k):
  - Given at time t(k-1): particle set (anywhere from 500 to 100 000)
  - Move each particle to next time using target motion model \Psi
  - Weight each particle using the "information update"
  - Resample and perturb → new particle set at time t(k)
- Resampling is crucial to particle filter success
  - At least one third of the particles are <u>not</u> resampled (!)
    → particle set collapse
  - Replicated particles are perturbed 'appropriately'
- Done carefully, particles are target trajectories over multiple time scans
  - Intuitive method: "Russian roulette and splitting"
    S. Ulam and J. von Neumann (1947)

iLRT Scoring

- iLRT = *intensity* likelihood ratio tracking (LRT)
- PHD / iFilter approximates the true multitarget posterior pdf
- Key idea: Use it as a "proposal density" to generate target track trajectories
- Score each trajectory using fastest possible statistical test for target present / target absent
  - Sequential LRT is optimal → uses fewest number of data scans
- Procedure
  - Step 1: Generate a collection of particle "paths".
    - Each path is a target trajectory
  - Step 2. Score each path using the iLRT

$$CLLR_k = CLLR_{k-1} + \ln \Psi_{k-1}(x_k \mid x_{k-1}) + \ln L_k(Z_k, U_k \mid x_k)$$

  - Step 3: Concatenate the CLLR to each path (a kind of junk DNA)
    - "Colorize" each path according to the CLLR

Example

- Scenarios are illustrative only
- All measurements and features are simulated
- Measurements z
  - Points contained in 75 x 75 km region
  - Errors are approximately 1 km in x and y, and are i. i. d.
- Feature u is assumed independent of target state x
  - H1: $p(u|z,x)=N(u|3,1)$
  - H0: $p(u|z,phi)=N(u|0,1)$
- Clutter
  - Uniform Poisson point process
  - Mean number of points is 10 or 20
- Time between pings = 180 s
- Target speeds approximately 8 knots
  - Target motion includes course and speed change models
- Probability of target detection is 0.1 or 0.2

Other Parameters

- Total particle count = 30 000
  - Total birthed particles = 10 000
    - Uniformly distributed
    - Total weight 0.02
  - No particle death
  - Resampling restores 30 000
- Memory requirements are ‘minimal’
  - X, Y, PHD weight, CLLR, etc… per particle
  - About 600MB for 1000 frames here (every particle kept)
  - 24 GB cost US$300 in 2011 CE
- Tracks are colored patriotically
  - They first turn “RED” when a threshold exceeded
    - Cumulative log likelihood ratio threshold = 20
  - Turn “WHITE” when a larger threshold exceeded
  - Then “BLUE”
- Display randomly selected 10% = 3000 particle tracks/ping
  - Causes weak tracks to flicker

Easy Multistatic Problem

- Two real targets
- Fantasy features "OFF"
- Pd = 0.2
- Average number of clutter data per scan is 10



Easy Multistatic Problem

- Two real targets
- Fantasy features "ON"
- Pd = 0.2
- Average number of clutter data per scan is 10



Easy Problem: Side-by-Side





Hard Multistatic Problem

- Three real targets
- Fantasy features "OFF"
- Pd = 0.1 (!!)
- Average number of clutter data per scan is 20



Hard Multistatic Problem

- Three real targets
- Fantasy features "ON"
- Pd = 0.1 (!!)
- Average number of clutter data per scan is 20



Hard: Side-by-Side





Track Extraction: Heritage Method

- Goal: To estimate from the particle set
  - Number of targets
  - For every target,
    - Point estimate of target state
    - "Error covariance" matrix
      - Issue: intensity filters do not yield error covariance matrices
- Key idea
  - Threshold the CLLR each particle
  - Optimum detector of target presence or absence along the path
    - Wald and Wolfawitz (c. 1955)
  - Trace the heritage of particles that are above threshold backward in time to their "origin" point
  - Distinct origins correspond to distinct targets
    - Targets are the particle clusters with a common origin
  - Extract point state estimates and covariance matrices from these target-specific clusters

Example With Strong Features

- One real target
- Four pingers, 400 pings
- For each ping, three source/receiver pairs
- Avg. number of clutter points per ping 47
- Coloring scheme is not quite "patriotic"
- After threshold exceeded
  - Find heritage points, or trajectory origins
  - Covariance matrices computed only after BLACK = All particles
  - BLUE = most "recent" particles



Multistatic Sonar Movies



Cluster Sharing

- If targets were in ‘close’ proximity at some past time, they will share particles



Cluster Sharing

- Problem can resolve itself by waiting



Egg Laying

- Maneuvering targets will often “lay an origin (egg)” at some time during the time of the maneuver

- Hybrid iLRT and iFilter
  - On-going work
  - To be reported at Fusion 2011 in Chicago
  - Further results in SDF in Berlin, October 2011

Single Target Bayes Filter

- Target state space $\mathbb{S}^+$ and measurement space $\mathbb{Z}$
- Target motion

$$\Psi_{X_k|X_{k-1}}(x \mid y) \text{ for all } x, y \in \mathbb{S}^+$$

- Measurement likelihood

$$p_{Z_k|X_k}(z \mid x) \text{ for all } z \in \mathbb{Z}, x \in \mathbb{S}^+$$

- Single-target Bayes filter
  - Prediction: $$p_{k|k-1}(x) = \int_{\mathbb{S}^+} \Psi_{X_k|X_{k-1}}(x \mid y)\, p_{k-1|k-1}(y)\, dy$$
  - Information update: $$p_{k|k}(x) = \frac{p_{Z_k|X_k}(z_k \mid x)\, p_{k|k-1}(x)}{p_{Z_k|Z_1,\ldots,Z_{k-1}}(z_k \mid z_1,\ldots,z_{k-1})}$$
- Multitarget Bayes filter
  - Unchanged, but target and measurement models are different

Multitarget Tracking

- Target state variables are PPPs
  - PPP event space is *huge* compared to single-target tracking
    - Simulation procedure (Lecture 1) shows how big it is
  - Bayes posterior pdf is defined but impractical to evaluate
- More sophisticated – and practical – alternative
  - PPPs are "parameterized" by their intensities
  - Use Bayes to update the posterior intensity function
- This agenda cannot be carried out exactly
  - Prior multitarget process is a PPP
  - Predicted multitarget process is a PPP
  - Predicted measurement process is a PPP
  - But the multitarget information updated point process is *not* a PPP
- What to do?
  - Approximate the information update with a PPP

Intensity Filter

- Bayes posterior pdf of the Multitarget information update
  - Explicit analytic form
  - Finite point process, but not a PPP
- Compute the single-target marginal pdfs
  - Exact analytic step
  - Observe that the marginal pdfs are identical
- Mean field approximation
  - Multitarget posterior pdf = product of marginal pdfs
- Define a PPP with intensity proportional to the marginal pdf
  - Scale it so the multitarget intensity "counts" targets (Streit and Stone, 2008)
  - Compare to Mahler's approach (c. 1995?)
  - -- Janossy density and the first moment approximation
  - <u>Mean field and first moment methods are essentially identical</u>

Predicted Multi-Target and Measurement PPPs

- The transformed target list is a realization of the predicted multitarget PPP on $\mathbb{S}^+$ with intensity

$$\boxed{f_{k|k-1}(x) = \int_{\mathbb{S}^+} f_{k-1|k-1}(y)\, \Psi_{X_k|X_{k-1}}(x|y)\, dy}$$

Proof. Intuitive counting argument

- The predicted measurement list is a realization of the predicted measurement PPP on $\mathbb{Z}$ with intensity

$$\boxed{\lambda_{k|k-1}(z) = \int_{\mathbb{S}^+} f_{k|k-1}(x)\, P_k^D(x)\, p_{Z_k|X_k}(z|x)\, dx}$$

Proof. Intuitive counting argument

Example of Measurement PDF

- Suppose there are $m = 2$ measurements: $\{z_1, z_2\}$
- The only realizations with nonzero probability have 2 states
- If both states are in S and $\neq \phi$, neither measurement is clutter:

$$p(z_1, z_2 \mid x_1, x_2) = p(z_1, z_2 \mid x_1, x_2)$$

$$= \frac{1}{2!}\left\{p(z_1 \mid x_1)\, p(z_2 \mid x_2) + p(z_1 \mid x_2)\, p(z_2 \mid x_1)\right\}$$

- If both states are $\{\phi, \phi\}$, then both measurements are clutter:

$$p(z_1, z_2 \mid \phi, \phi) = \frac{1}{2!}\left\{p(z_1 \mid \phi)\, p(z_2 \mid \phi) + p(z_1 \mid \phi)\, p(z_2 \mid \phi)\right\}$$

- If one state is ϕ, exactly one of the measurements is clutter:

$$p(z_1, z_2 \mid x_1, \phi) = p(z_1, z_2 \mid \phi, x_1)$$

$$= \frac{1}{2!}\left\{p(z_1 \mid x_1)\, p(z_2 \mid \phi) + p(z_1 \mid \phi)\, p(z_2 \mid x_1)\right\}$$

Multitarget Measurement PDF

- $\xi = (m, x_1, \ldots, x_m)$ = realization of predicted multitarget PPP on S^+
- Measurement set: $\upsilon = (r, z_1, \ldots, z_r)$ need <u>not</u> be a PPP
- All data arise from targets in the augmented state space S^+
  - One target per measurement, and conversely
    - If target state is ϕ, then measurement is clutter
    - If target state is x, then measurement is from target with state x
  - Target assignments are unknown
  - All assignments equally likely *a priori*
- Multitarget likelihood function is sum over all permutations

$$p_{\Upsilon_k \mid \Xi_k}(\upsilon \mid \xi) = \begin{cases} \dfrac{1}{m!} \displaystyle\sum_{\sigma \in Sym(m)} \prod_{j=1}^{m} p_{Z_k \mid X_k}\left(z_j \mid x_{\sigma(j)}\right) & \text{if } r = m \\ 0 & \text{if } r \neq m \end{cases}$$

- Only target realizations with m points are nonzero probability
- Augmented state space avoids explicitly enumerating clutter

Exact Multitarget Information Update

- Bayes information updated pdf

$$\boxed{p_{k|k}(x_1, \ldots, x_m)} \equiv p_{\Upsilon_k|\Xi_k}(\upsilon|\xi) \frac{p_{\Xi_k|\Upsilon_1,\ldots,\Upsilon_{k-1}}(\xi)}{p_{\Upsilon_k|\Upsilon_1,\ldots,\Upsilon_{k-1}}(\upsilon)}$$

- Exactly the same form as single-target update
- Substitute predicted target & measurement PPP likelihood functions

$$p_{k|k}(x_1, \ldots, x_m) = \frac{1}{m!} \sum_{\sigma \in Sym(m)} \prod_{j=1}^{m} \frac{p_{Z_k|X_k}\left(z_j \mid x_{\sigma(j)}\right) P_k^D\left(x_{\sigma(j)}\right) f_{k|k-1}\left(x_{\sigma(j)}\right)}{\lambda_{k|k-1}\left(z_j\right)}$$

- Posterior target point process is NOT a PPP because:

$$p_{k|k}(x_1, \ldots, x_m) = 0$$

if: # measurements $\neq$ # target states (including ϕ)

Mean Field Approximation

$$p_{k|k}(x_1,\ldots,x_m)=\frac{1}{m!}\sum_{\sigma\in Sym(m)}\prod_{j=1}^{m}\frac{p_{Z_k|X_k}\left(z_j\mid x_{\sigma(j)}\right)P_k^D\left(x_{\sigma(j)}\right)f_{k|k-1}\left(x_{\sigma(j)}\right)}{\lambda_{k|k-1}\left(z_j\right)}$$

- Denominators are the integrals of the numerators since, for any z,

$$\lambda_{k|k-1}(z)=\int_{S^+}p_{Z_k|X_k}(z\mid x)P_k^D(x)\,f_{k|k-1}(x)\,dx$$

- The integral (marginal) of the joint pdf ovet all targets ... except x_s

$$p_{k|k}(x_s)=\int_{S^+}\cdots\int_{S^+}p_{k|k}(x_1,\ldots,x_m)\underbrace{dx_1\ldots dx_{s-1}dx_{s+1}\ldots dx_m}_{omit\ dx_s}$$

$$=\frac{1}{m}\sum_{j=1}^{m}\frac{p_{Z_k|X_k}\left(z_j\mid x_s\right)P_k^D\left(x_s\right)f_{k|k-1}\left(x_s\right)}{\lambda_{k|k-1}\left(z_j\right)}$$

$$=\frac{1}{m}f_{k|k}^D\left(x_s\right)$$

- Important: Marginal pdfs are identical for all s

Mean Field Approximation

- Product of the marginal pdfs is a symmetric function

$$p_{k|k}(x_1, \ldots, x_m) \cong \prod_{s=1}^{m} p_{k|k}(x_s)$$

- Scaling is needed to enable the intensity filter to "count" targets
- Determine scale factor by MAP estimation

$$\frac{1}{m!}\exp\left(-\int c\, p_{k|k}(x)dx\right)\prod_{j=1}^{m}\left(c\; p_{k|k}(x_j)\right) \propto \exp(-c)\; c^m$$

$$\Rightarrow \hat{c}_{MAP} = m$$

- Intensity of the approximating PPP for the detected targets

$$f_{k|k}^{D}(x) = \sum_{j=1}^{m} \frac{p_{Z_k|X_k}(z_j \mid x)\, P_k^D(x)\, f_{k|k-1}(x)}{\lambda_{k|k-1}(z_j)}$$

- PPP approximation closes the PHD / intensity filter recursions

How To Speak Janossy

- Janossy density is defined for any finite point process

$$j_n\left(x_1,\ldots,x_n\right)=\Pr\begin{bmatrix}\text{Exactly } n \text{ points in a realization}\\ \text{with one point in each infinitesimal}\\ [x_i,\ x_i+dx] \text{ for } i=1,\ \ldots,\ n\end{bmatrix}$$

- The Bayes posterior target process is a finite point process

$\boxed{\textit{for } n=m:}$

$$j_m(x_1,\ldots,\ x_m)=m!\,p_{k|k}(x_1,\ \ldots,\ x_m)$$

$\boxed{\textit{for } n\neq m:}$ $\quad j_n(x_1,\ldots,\ x_n)=0$

- In PHD and intensity filters, only one Janossy density is nonzero (!)

How to Speak in First Moments

- Definition (Daley and Vere-Jones, First Edition 1988)

$$\mathrm{M}_1(x) \equiv \sum_{n=0}^{\infty} \frac{1}{n!} \int_{S^+} \cdots \int_{S^+} j_{n+1}\left(x, \boxed{x_1, \ldots, x_n}\right) \boxed{dx_1 \ldots dx_n}$$

  – No intuitive interpretation seems available in the literature

- For Bayes posterior target process, only term $n+1=m$ is nonzero, so

$$\mathrm{M}_1(x) = \frac{1}{(m-1)!} \int_{S^+} \cdots \int_{S^+} j_m\left(x, x_1, \ldots, x_{m-1}\right) dx_1 \ldots dx_{m-1}$$

- Substitute and simplify

$$\mathrm{M}_1(x) = \frac{m!}{(m-1)!} \int_{S^+} \cdots \int_{S^+} p_{k|k}\left(x, x_1, \ldots, x_{m-1}\right) dx_1 \ldots dx_{m-1}$$

$$= \frac{m!}{(m-1)!} \frac{1}{m} f_{k|k}^{D}(x)$$

$$= f_{k|k}^{D}(x)$$

- Scaled mean field approximation = first moment approximation

Regularization

Regularization

- PET imaging is *stochastic* inverse problem: data are noisy
  - Intrinsic image reconstruction artifacts akin to Gibbs oscillations
  - Noise artifact worsens as number of Shepp-Vardi iterations increase (!)
- Necessity for regularization in PET is not surprising
  - Number of data is finite
  - High resolution and one intensity parameter per pixel spell "*Trouble*"
- Can PHD and Intensity filters escape the problem?
- Regularization methods for PET
  - Markov random fields, Penalty functions, Resolution kernels *
  - Grenander's Method of Sieves *

    Sieve restricts intensity to a smooth subspace

    $$\left\{ \lambda(x) = \int_U k(x \mid u)\ \varsigma(u)\ du \ \middle|\ \varsigma(u) \geq 0 \right\}$$ where U is any convenient space

    and kernel $k(\cdot)$ is specified pdf $\int_R k(x \mid u)\ dx = 1$ for all $u \in U$

    -- In other words, the kernel "low pass filters" the intensity $\varsigma(u)$
  - Ridge following methods

* D. L. Snyder and M. I. Miller, Random Point Processes in Time and Space, 1991

Cramer-Rao Bound (CRB) For Unbiased Intensity Estimators

General CRB

- Let θ denote the parameter vector of a PPP intensity $\lambda(x;\theta)$
  - -- Assume first two derivatives of $\lambda(x;\theta)$ with respect to θ exist
- CRB for unbiased estimators of θ using a realization of the PPP is inverse of

$$FIM(\theta) = E\left[\left(\nabla_\theta \log p(\xi;\theta)\right)\left(\nabla_\theta \log p(\xi;\theta)\right)^T\right]$$ Fisher Information Matrix

$$= \int_R \frac{1}{\lambda(s;\theta)} \left(\nabla_\theta \lambda(s;\theta)\right)\left(\nabla_\theta \lambda(s;\theta)\right)^T ds$$ = An ordinary integral!

  - -- where the expectation is with respect to the PPP
- Loglikelihood of the PPP realization $\xi = \left(n, \{x_1, \ldots, x_n\}\right)$ is

$$\log p(\xi;\theta) = -\int_R \lambda(s;\theta)\, ds + \sum_{j=1}^{n} \log \lambda\left(x_j;\theta\right)$$

$$\Rightarrow \quad \nabla_\theta \log p(\xi;\theta) = -\int_R \nabla_\theta \lambda(s;\theta)\, ds + \sum_{j=1}^{n} \frac{1}{\lambda\left(x_j;\theta\right)} \nabla_\theta \lambda\left(x_j;\theta\right)$$

- Expectation of a random size sum: Campbell's Theorem (1909). QED

CRB for PET Imaging

- Recall PET measurement intensity

$$\mu(y;\Lambda) = \int_{\mathrm{R}} p(y \mid x)\, \lambda(x)\, dx = \sum_{k=1}^{K} \lambda_k\, f_k(y)$$

where $\Lambda = \left(\lambda_1, \ldots, \lambda_K\right)$

$$f_k(y) = \int_{R_k} p(y \mid x)\, dx \qquad \{R_k\} = \text{pixels of image } \mathrm{R}$$

$$\Rightarrow \quad \nabla_\Lambda\, \mu(y;\Lambda) = \left[\, f_1(y) \quad \cdots \quad f_K(y) \,\right]^T \in \mathbb{R}^K$$

- For PET: $$FIM\left(\Lambda\right) = \left[\int_{\mathrm{R}} \frac{f_j(y)\, f_k(y)}{\mu(y;\Lambda)}\, dy \right]_{j,k=1}^{K}$$

- Computing $CRB = FIM^{-1}$ is impractical for high resolution, i.e., large K
- Extensions to PHD / Intensity filters are currently unavailable
- Hero et al. (1992, 1994) give fast $O\left(K\right)$ time iterative method
  - -- Compute CRB of pixels in "small" sub-regions of interest
  - -- *Small* regions could be targets (or tumors)

Comparisons to

Single Target

Likelihood Ratio Detection and Tracking (LRT)

Comparison of LRDT Update Formula to PHD and Multitarget iFilter*

LRDT: $$L_{k|k}(s) = L_{k|k-1}(s)\left[1 - P_d(s) + \sum_{i=1}^{n_k} \frac{P_d(s)p(y_i|s)}{\lambda w(y_i)}\right]$$

iFilter: $$f_{k|k}(s) = f_{k|k-1}(s)\left[1 - P_d(s) + \sum_{i=1}^{n_k} \frac{P_d(s)p(y_i|s)}{\lambda^{Clutter}_{k/k-1}(y_i) + \int_S P_d(s)p(y_i|s)f_{k|k-1}(s)ds}\right]$$

PHD: $$f_{k|k}(s) = f_{k|k-1}(s)\left[1 - P_d(s) + \sum_{i=1}^{n_k} \frac{P_d(s)p(y_i|s)}{\lambda(y_i) + \int_S P_d(s)p(y_i|s)f_{k|k-1}(s)ds}\right]$$

- **The missing term in denominator of the LRDT reflects the assumption of 0 or 1 targets present.**
- **The λ term in all three represents the false target or clutter intensity. In LRDT and PHD, this is an assumed known prior intensity. In the iFilter, clutter intensity is estimated recursively.**

**Slide adapted from Plenary Presentation, "Likelihood Function Detection and Tracking," L. Stone, ISIF International Conference on Information Fusion, Seattle, July 2009.*

Single Photon Emission Computed Tomography (SPECT)

SPECT

- Low resolution compared to PET
  - Older than PET – dates to the mid-1960s
  - SPECT procedure costs $500; PET costs $30000 (approx., 2008)
    - Much more widely used than PET
- SPECT is inherently three dimensional, unlike PET
- Radioisotope given to patient
  - Decays into gamma photons
  - Two dimensional "snapshots" taken by a gamma camera
    - Also called an Anger camera after its developer, H. Anger, 1957
  - Camera is moved to various positions and a snapshot taken in each position
- Three dimensional image reconstructed from snapshots
  - Miller, Snyder, and Miller, IEEE Trans. Nuclear Science, 1985

Gamma Cameras



Gamma Camera in One Position



Required Functions

- Sensor measurement model

$$p_{Y|X\Theta}(y \mid x, \theta) = pdf\left(\begin{array}{c|c} \text{measured photon location at} & \text{decay location at 3-D} \\ \text{2-D point } y \text{ on camera face} & \text{point } x \text{ and view angle } \theta \end{array}\right)$$

  -- Models the quality of the sensor at location θ

  -- Called the blurring function in image restoration apps

- Environmental propagation model

$$\beta(x, y, \theta) \;=\; \Pr\left[\begin{array}{c} \text{photon emitted at 3-D decay point } x \text{ moving} \\ \text{toward location } y \text{ on the camera face at} \\ \text{view angle } \theta \text{ arrives at the camera face} \end{array}\right]$$

  – Depends on the camera geometry and the attenuation

- Sensor opportunity

$$p_{\Theta|X}\left(\theta_j \mid x\right) = \Pr\left[\begin{array}{c} \text{photon emanating from } x \text{ propagates} \\ \text{toward the camera face with view angle } \theta \end{array}\right]$$

  – Models, e.g., target aspect dependence

E-Step of EM: Auxiliary Function

- Missing data correspond to
  - Random number N_j of missed photon detections in camera at view angle θ_j
  - Locations $\left\{ y_{jr} \right\}_{r=1}^{N+n_j}$ of undetected photons
- Auxiliary function. EM iteration index = m

$$Q = \int_{\square^3} \left\{ -\lambda(x) + \left[\left(1 - \bar{\beta}(x)\right) \lambda^{(m)}(x) + f^{(m)}(x) \right] \log \lambda(x) \right\} \mathrm{d}x$$

where $$\bar{\beta}(x) = \sum_{j=1}^{n_\Theta - 1} p_{\Theta|X}\left(\theta_j \mid x\right) \beta_j(x)$$

$$\beta_j(x) = \int_{Face(\theta_j)} p_{Y|X\Theta}\left(y \mid x, \theta_j\right) \beta\left(x, y, \theta_j\right) \mathrm{d}y$$

$$f^{(m)}(x) = \sum_{j=1}^{n_\Theta - 1} \sum_{r=1}^{n_j} w_j\left(x \mid y_{jr}, \theta_j\right)$$

$$w_j\left(x \mid y_{jr}, \theta_j\right) = \frac{p_{Y|X\Theta}\left(y \mid x, \theta_j\right) p_{\Theta|X}\left(\theta_j \mid x\right) \beta\left(x, y, \theta_j\right) \lambda^{(m)}(x)}{\int_{\square^3} \quad " \quad \mathrm{d}x}$$

M-Step

- Use calculus of variations to max Q over choice of $\lambda(x)$
- SPECT recursion:

$$\lambda^{(m+1)}(x) = \lambda^{(m)}(x) \sum_{All\ cameras\ \{\theta_j\}} p_{\Theta|X}\left(\theta_j \mid x\right)$$

$$\times \underbrace{\left[1 - \beta_j(x) + \sum_{\substack{data\ \{y_{jr}\} \\ camera\ j}} \frac{p_{Y|X\Theta}\left(y_{jr} \mid x, \theta_j\right) \beta(x, y_{jr}, \theta_j)}{\int_{\square 3} p_{Y|X\Theta}\left(y_{jr} \mid x, \theta_j\right) p_{\Theta|X}\left(\theta_j \mid x\right) \beta(x, y_{jr}, \theta_j) \lambda^{(m)}(x)\, dx} \right]}_{\equiv\ L_j(x)}$$

- Weighted average of single-camera intensity estimates

$$\boxed{\lambda^{(m+1)}(x) = \lambda^{(m)}(x) \sum_{All\ cameras\ \{\theta_j\}} p_{\Theta|X}\left(\theta_j \mid x\right) L_j(x)}$$

Multisensor Intensity Filter

- Multisensor PHD and intensity filter assumes PPP target model
- Sensors are therefore independent, not conditionally independent
- The PPP intensity is the average

$$f_{k|k}^{\text{Fused}}(x) = \frac{1}{M}\left(\sum_{\ell=1}^{M} L_k^{\ell}\left(x \mid Z_k\right)\right) f_{k|k-1}^{\text{Fused}}(x)$$

where M is the number of sensors, possibly different

Z_k is the measurement set at time t_k

and the sensor-level intensity filter information update is

$$L_k^{\ell}\left(x \mid Z_k\right) = 1 - P_k^D(x;\ell) + \sum_{j=1}^{m} \frac{p_k^{\ell}(z_j \mid x)\, P_k^D(x;\ell)}{\int_{S^+} p_k^{\ell}(z_j \mid s)\, P_k^D(s;\ell)\, f_{k|k-1}^{\text{Fused}}(s)\, ds}$$

- Proof. Use either mean field and MAP scaling, or first moment

Variance Reduction

Estimated Number of Targets

- For identical sensor coverages

$$N_{k|k}^{\text{Fused}} = \int_S f_{k|k}^{\text{Fused}} \, \mathrm{d}x$$

$$= \int_S \frac{1}{M}\left(\sum_{\ell=1}^{M} L_k^{\ell}\left(x \mid Z_k\right)\right) f_{k|k-1}^{\text{Fused}}(x) \, \mathrm{d}x$$

$$= \frac{1}{M}\sum_{\ell=1}^{M} N_{k|k}^{\ell}$$

- Estimates are Poisson distributed $\therefore$ variance is equal to the mean
  - Assuming $N_{k|k}^{\ell}$ is unbiased, then

$$Var\left[N_{k|k}^{\text{Fused}}\right] = \frac{1}{M}\sum_{\ell=1}^{M} Var\left[N_{k|k}^{\ell}\right]$$

$$= \frac{1}{M}\sum_{\ell=1}^{M} Mean\left[N_{k|k}^{\ell}\right] = \frac{\bar{N}_{true}}{M}$$

- The more sensors the better

Concluding Remarks

- Intensity and PHD filters are *tomographic* imaging filters
  - Image space is target state space
  - Based on flawed physical model of real targets
  - Deserve careful investigation
- Intensity filters linked to important signal processing problems
  - PET for single sensors
  - SPECT for multiple sensors
- These links are important
  - Physical interpretation of the multiple target model
  - Convincing intuitive understanding of the information updates
- Important concepts need further study
  - Cramer-Rao Bound for the intensity function
    - CRB for PET on a finite grid is presented
    - Needed: PCRB (Posterior Cramer-Rao Bound)
  - Regularization
    - Reconstruction artifacts widely recognized phenomena in tomography
  - Estimates of target birth and clutter PPPs
    - Uses the augmented target state space S+
  - Variance reduction via multi-sensor averaging

Poisson Point Processes: Imaging, Tracking, Sensing

- R. Streit, Springer, New York, 2010

Chapters / Table of Contents

Part I: Background

- Basic properties
- Intensity Estimation
- Cramer-Rao Bounds for Intensity Estimates

Part II: Applications

- Tomographic Imaging
- Multitarget Tracking
- Distributed Sensing

Part III: Beyond the Poisson Point Process